

PiMPeR: Piecewise Dense 3D Reconstruction from Multi-View and Multi-Illumination Images

Reza Sabzevari¹, Vittori Murino², and Alessio Del Bue²

¹ Robotics and Perception Group, University of Zurich, Switzerland

² Pattern Analysis and Computer Vision (PAVIS),
Italian Institute of Technology, Genova, Italy

March 18, 2015

Abstract

In this paper, we address the problem of dense 3D reconstruction from multiple view images subject to strong lighting variations. In this regard, a new piecewise framework is proposed to explicitly take into account the change of illumination across several wide-baseline images. Unlike multi-view stereo and multi-view photometric stereo methods, this pipeline deals with wide-baseline images that are uncalibrated, in terms of both camera parameters and lighting conditions. Such a scenario is meant to avoid use of any specific imaging setup and provide a tool for normal users without any expertise. To the best of our knowledge, this paper presents the first work that deals with such unconstrained setting. We propose a coarse-to-fine approach, in which a coarse mesh is first created using a set of geometric constraints and, then, fine details are recovered by exploiting photometric properties of the scene. Augmenting the fine details on the coarse mesh is done via a final optimization step. Note that the method does not provide a generic solution for multi-view photometric stereo problem but it relaxes several common assumptions of this problem. The approach scales very well in size given its piecewise nature, dealing with large scale optimization and with severe missing data. Experiments on a benchmark dataset *Robot data-set* show the method performance against 3D ground truth.

1 Introduction

This paper addresses the problem of dense reconstruction from several uncalibrated multi-view images that are subject to lighting variation. Such a dense reconstruction pipeline is meant to deal with a off-the-shelf hardware setup—both in terms of image acquisition and processing—in a reasonable amount of

time. The proposed method is based on two fundamental reconstruction methods and accordingly consists of two main stages: first, using *Structure from Motion* (SfM) to recover the coarse geometry of the scene/shape, and second using *Photometric Stereo* (PS) to recover fine surface details exploiting photometric properties of the surface. In this way, the method can take maximum advantage of geometric and photometric scene properties captured by multi-view and multi-illumination images. For this reason, we call our proposed pipeline as PiMPeR stands for *Piecewise Multi-view Photo-geometric dense Reconstruction*. A sparse 3D point cloud is obtained using SfM by tracking and matching a set of sparse feature points on the image sequence. This step also provides the camera positions, which will be used in the further steps of the pipeline. Such a 3D point cloud will be too sparse, because very few feature tracks will survive among the images under drastic variation of lighting conditions. The coarse 3D mesh is defined by first projecting the point cloud on one of the images and performing 2D Delaunay triangulation, and then applying the same triangulation on the point cloud. Such a 3D mesh allows us to decompose the images into a set of image patches to simplify the establishment of multi-view correspondences and it further helps to assemble the fine-detailed surface patches to form a dense surface. Once the multi-view images are decomposed, the corresponding image patches are registered on a template patch and the PS is used to recover dense surfaces for all the patches. This requires adoption of classic PS methods and their reformulation in a piecewise framework, which consequently requires an integration step to collect all surface patches to a uniform one.

1.1 Comparisons with related work

Recent advancements in rigid SfM have pushed the boundaries in different aspects, such as large scale 3D modelling in urban areas [1], real-time dense 3D reconstruction [2] and, most recently, chronological scene reconstruction [3]. These methods have either assumptions on fixed lighting conditions in real-time scenarios or use a tremendous amount of images to find a much smaller number of points that can be tracked across the images. However, in realistic imaging conditions, the scene illumination is usually subject to change. Moreover, the variation of the lighting conditions is fundamental to accurately reconstruct parts of the scene with homogeneous texture. Recently, Pizzoli et al. [4] and Engel et al. [5] showed that only high gradient textures could be used as reliable observations of the scene for 3D reconstruction, and that, to recover 3D position of the scene parts with homogeneous textures, strong prior information such as surface smoothness is essential.

Differently, we present a novel method that explicitly takes into account changing illumination and uses them to obtain dense 3D reconstructions in an uncalibrated and multi-view setup. In order to tackle this problem, the classic formulation of PS is adapted in a piecewise framework on multi-view image patches registered on a template patch with a considerable amount of missing data. The dense matching of multi-view image patches is possible by multi-view constraints imposed by a set of sparse image points that could be reliably

tracked over the images.

To this end, our approach proposes a multi-view piecewise decomposition of the images that permits us to solve locally both the image pixel correspondence and the photometric stereo problems. In particular, we are able to deal with the most difficult case of wide-baseline images, where given both lighting and geometric distortions, we have image patches with strong scale changes, self-occlusion, cast-shadows, and strong specularities. Previous approaches, instead strongly rely on small baselines or custom setups in order to establish reliable correspondences between pixels in different frames [6, 7, 8, 9]. Other geometric methods that do not exploit illumination variation rely on large datasets of images and require powerful processing resources [10]. In such a scenario, our proposed framework is a step forward towards dense reconstruction from a few uncalibrated images with arbitrary camera motion and lighting conditions.

Factorization-based methods for PS have certainly provided efficient closed-form solutions without using specific equipment such as structured light. It is based on the fact that a set of images taken from a static view and subject to varying lights lies in a certain subspace. Hayakawa [11] first made evident such constraints assuming a Lambertian surface and a single light source. Basri et al. [12] used a more descriptive photometric model based on a spherical harmonic representation of lighting variations. These classic methods in PS, which do not have any depth assumption, always tie with the bas-relief ambiguity [13]. On the contrary, Shi et al. [14] perform an automatic radiometric calibration by identifying a new set of constraints that can solve for the Generalised Bas-Relief (GBR) transformation. Finally, Papadimitri and Favaro [15] proposed a closed-form solution based on the identification of a set of local diffuse reflectance maxima which provides enough constraints to solve for the GBR ambiguity.

Another viable option to obtain reliable and detailed 3D surfaces is to use a structured pattern for illuminating the object in a controlled environment [16]. These active systems have led to a high number of custom solutions [17] that require a laboratory setup and accurate calibration of the devices. Hernandez et al. [18], instead, use a less restrictive calibrated setup with three not collinear colored lights in a dark room with surfaces that also need to be photometrically calibrated. Recently, Anderson et al. have extended this approach to arbitrary colored surfaces [19]. Finally, with a less constrained setup, the method of Park et al. [20] considers a coarse-to-fine approach that requires only a stereo pair.

More related to our approach, Multi-View Photometric Stereo (MVPS) use both geometric and photometric constraints to solve the problem in multi-view setups but using short baselines [6]. Similarly, Lim et al. [7] use multiple views to generate a coarse planar surface of the object as a 3D triangulated mesh based on recovered sparse 3D points. Then, they perform PS in an iterative way for each triangle to recover the dense surface and to align it with the recovered 3D points. They also assume the lighting conditions to be fixed and that they can reconstruct the objects having small baselines only. Joshi and Kriegman [8] define the multi-view photometric algorithm as the optimization of multi-view and photometric matching costs using a graph-cut based approach. The main drawback of the method is an inherent fronto-parallel assumption that leads to

local rank constraints in the images.

In general, a piecewise formulation helps to deal with changes of both geometrical and photometric components over time by simplifying the multi-view correspondence problem. Piecewise methods have recently demonstrated their strength for solving SfM problems of increasing complexity (see [21], [22] and [23]).

1.2 Contribution

Considering the problem that is addressed in this paper, one of the most challenging scenarios for dense 3D reconstruction has been targeted. We do not use a specific imaging setup and we consider images from uncalibrated cameras with arbitrary lighting conditions and view-points. Unlike other methods (i.e. [5, 10]), that require many images and huge amounts of processing resources to create fairly dense point clouds, we use a few images and the processing capacities available on normal PCs to obtain dense surfaces.

Differently from [8, 7, 6], we deal with a wide-baseline scenario where the images can be taken after consistent camera motion, which is the current weakness of MVPS algorithms. Using a piecewise formulation allows us to solve for the MVPS problem using local estimates of the photometric properties of the shape, which results in efficiently by solving several small optimization problems rather than a global one. Using piecewise photometric stereo (PPS) we can deal with local shape ambiguities using our local-to-global surface patches alignment.

1.3 Structure of the Paper

The rest of the paper is structured as follows. Section 2 describes PiMPeR, our pipeline for photo-geometric dense reconstruction. It details all the steps of our pipeline: the generation of a coarse mesh, the recovery of detailed surface patches and the specification of the optimization problem that is used to assemble the detailed surface patches on the coarse mesh. Experiments in Section 3 present the 3D reconstruction obtained by our approach and, finally, Section 4 draws the conclusions and the path for future research potentials.

2 The PiMPeR Pipeline

Our approach combines both multi-view and photometric constraints to obtain a complete 3D reconstruction of the object together with its photometric properties. The complete pipeline is sketched in Figure 1 and we briefly resume each stage before describing them in more detail in the next sections. Note that the input to our algorithm is only a set of uncalibrated images, object albedo and lighting directions are not given as well. First we create a coarse 3D mesh using SfM by tracking/matching a set of feature points in the image sequence.

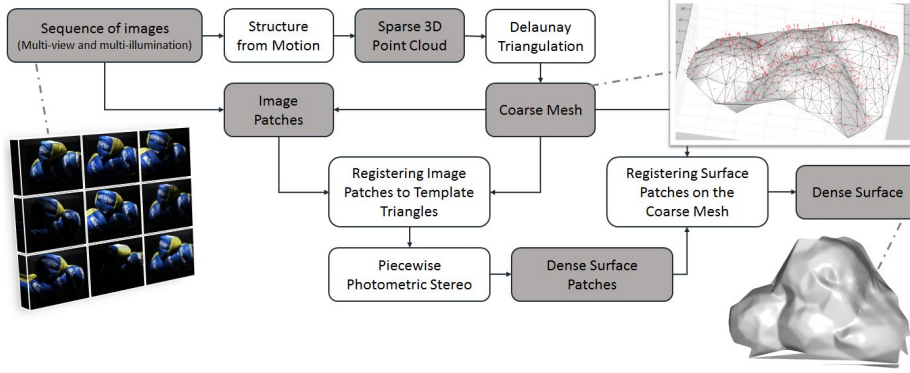


Figure 1: The PiMPeR pipeline: input sequence contains images with different viewpoints and arbitrary lighting; the coarse 3D mesh generated by SfM and then projected onto the images; the input sequence is decomposed to triangular image patches; triangular image patches are registered on a template triangle, which is the corresponding facet on the coarse mesh; detailed surface patches are recovered using piecewise photometric stereo; the complete surface is obtained by globally registering all the surface patches on the coarse mesh.

Given the adverse lighting conditions, we account for large percentages of missing data in the tracked 2D features by adopting robust SfM methods. Next stages. Each triangle of the mesh corresponds to a set of image patches. These image patches are first registered to a common image template and then the corresponding surface patches are reconstructed using a local photometric stereo method. At the final stage, the geometric constraints on the surface—defined by the 3D mesh—are used to assemble the surface patches, taking advantage of the proposed photo-geometric alignment.

2.1 Mesh Generation

The PPS reconstruction requires an initial partition of the images into a set of patches. For this reason, a 3D mesh is computed by performing multi-view 3D reconstruction from a set of 2D coordinates tracked/matched in the image sequence. This task can be efficiently done using a robust method for SfM that can cope with a large number of missing data in the measurements followed by a 3D mesh generation stage (see [24]). In more detail, after recovering the 3D positions belonging to p image correspondences via a rigid SfM and storing them in matrix \mathbf{S}_{sfm} , a Delaunay triangulation is computed to generate a coarse 3D mesh having triangular image patches. Such a set \mathcal{T} of m triangles is given by:

$$\mathcal{T} = \{\mathbf{t}_1 \dots \mathbf{t}_m\}, \quad \mathbf{t}_\mu \in \{1 \dots p\}^3 \mid \mu = 1 \dots m \quad (1)$$

where \mathbf{t}_μ contains the 3 indices identifying the 3 vertices of the triangle μ . Now, the vertices of triangle μ for each image frame g , ($g = 1 \dots f$), form a 2×3

matrix $\mathbf{W}_{g\mu}$. Likewise, the 3D triangle vertices are stored in a 3×3 matrix \mathbf{Z}_μ . This defines two sets of matrices that form the 3D mesh as $\mathcal{S} = \{\mathbf{Z}_\mu\}$ and the 2D mesh at each frame g as $\mathcal{W}_g = \{\mathbf{W}_{g\mu}\}$.

In other words, the triangulation is done on 2D image points of an arbitrary view (i.e., the frontal view) and then the resulted point indices that define the 2D triangulated mesh are used to triangulate the 3D points.

This mesh is then used for three purposes: decomposing the multi-view image frames into triangular image patches, assembling the recovered triangular surface patches in global 3D shape, and correcting the local depth ambiguities.

2.2 Multi-View Image Patches Decomposition

Once the mesh is generated, we proceed on decomposition by partitioning the image sequences into pieces that will be reconstructed using a local photometric stereo algorithm.

Consider an image as a matrix \mathbf{I}_g of size $h \times w$. To perform the decomposition, the 3D mesh \mathcal{S} is projected onto the image frames $\{\mathbf{I}_g\}$, $g \in [1 \dots f]$, thus, providing $m \times f$ triangles. A triangular patch is represented as a set of $h \times w$ binary matrices $\mathbf{B}_{g\mu}$. Using such binary masks, pixels included in the triangle μ can be extracted from each image frame. In order to ease the following steps of 3D patch alignment (Sec. 2.5), we enlarge the binary masks having a set of overlapping pixels with their immediate neighbouring patches.

Let $\mathbf{W}_{g\mu}$ be the projection of triangle t_μ on image frame g and consider the enlargement of triangle t_μ such that the new 2D vertices on the image plane define a bigger triangle in which the new edges are parallel to the original ones. This enlarged triangle is defined as \tilde{t}_μ with 2D vertices denoted by $\tilde{\mathbf{W}}_{g\mu}$. The conversion from 2D vertices $\tilde{\mathbf{W}}_{g\mu}$ to an image binary mask is made using the three barycentric coordinates $\alpha_{gk}^{(\mu)}$, $\beta_{gk}^{(\mu)}$ and $\gamma_{gk}^{(\mu)}$. Under this coordinate system, a pixel k that belongs to triangle \tilde{t}_μ projected on frame \mathbf{I}_g must satisfy the following condition:

$$\alpha_{gk}^{(\mu)} + \beta_{gk}^{(\mu)} + \gamma_{gk}^{(\mu)} = 1, \quad (2)$$

where

$$k \in \{1, \dots, b\}, \quad \mathcal{K}_{g\mu} = \left\{ k : \alpha_{gk}^{(\mu)} + \beta_{gk}^{(\mu)} + \gamma_{gk}^{(\mu)} = 1 \right\} \quad (3)$$

with $\mathcal{K}_{g\mu}$ being the set of indices for pixels belonging to triangle \tilde{t}_μ on frame g . Using this formalization, the binary mask for the projection of triangle \tilde{t}_μ on frame g is defined as a $h \times w$ matrix:

$$\mathbf{B}_{g\mu} = \begin{bmatrix} B_{g1}^{(\mu)} & \dots & B_{g(b-h)}^{(\mu)} \\ & \ddots & \\ B_{gh}^{(\mu)} & \dots & B_{gb}^{(\mu)} \end{bmatrix}_{h \times w}, \quad B_{gi}^{(\mu)} = \begin{cases} 0 & i \notin \mathcal{K}_{g\mu} \\ 1 & i \in \mathcal{K}_{g\mu} \end{cases}. \quad (4)$$

So, every image patch is given as:

$$\mathbf{J}_{g\mu} = \mathbf{B}_{g\mu} \odot \mathbf{I}_g, \quad (5)$$

where operator \odot stands for the element-wise multiplication.

2.3 Multi-View Image Pixel Registration

Each triangle on the coarse 3D mesh corresponds to a set of image patches that are obtained by projecting that particular triangle on the images. Given the camera (or object) motion, a set of image patches with different shapes and sizes will be obtained for every triangular mesh facet. This fact introduces a challenge, because pixel-to-pixel image patch correspondence has to be solved. Moreover, since the patches may consistently vary their sizes, some patches might have missing pixels with respect to their corresponding reference patch.

For this reason we need a registration process that aligns all the corresponding image patches in different views. To that end, for each triangle \tilde{t}_μ we define a template and register all image patches to their corresponding templates. Consider $\tilde{\mathbf{Z}}_\mu$ as a 3×3 matrix including the 3D coordinates of enlarged triangle \tilde{t}_μ from the coarse 3D mesh \mathcal{S} . Such 3D triangles have a normal \mathbf{N}_f^μ aligned to be fronto-parallel to the image plane [25] (i.e., aligned to $\mathbf{N}_I = [0 \ 0 \ 1]^\top$). The template triangle $\bar{\mathbf{Q}}_\mu$ is obtained by estimating a rotation matrix \mathbf{R}_μ that transforms the corresponding mesh facet and then projects the rotated triangle on an arbitrary image plane, such that:

$$\bar{\mathbf{Q}}_\mu = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \cdot \mathbf{R}_\mu \cdot \tilde{\mathbf{Z}}_\mu. \quad (6)$$

Hence, $\bar{\mathbf{Q}}_\mu$ is a 2×3 matrix containing the 3 coordinates of the template triangle as:

$$\bar{\mathbf{Q}}_\mu = [\bar{\mathbf{q}}_{\mu 1} \ \bar{\mathbf{q}}_{\mu 2} \ \bar{\mathbf{q}}_{\mu 3}]. \quad (7)$$

A practical solution for registering multi-view triangular image patches to the corresponding templates is given by the barycentric coordinate system. Consider pixel k from image frame \mathbf{I}_g belonging to triangle \tilde{t}_μ . This pixel coordinate can be determined as a 3-element barycentric coordinate:

$$\boldsymbol{\lambda}_{gk}^{(\mu)} = \begin{bmatrix} \alpha_{gk}^{(\mu)} & \beta_{gk}^{(\mu)} & \gamma_{gk}^{(\mu)} \end{bmatrix}^\top. \quad (8)$$

In this way, every pixel k belonging to triangle \tilde{t}_μ in all image frames can be mapped to a single position (with respect to barycentric coordinates) in the template triangle given by $\bar{\mathbf{Q}}_\mu$. This piecewise planar mapping gives the pixel-to-pixel correspondences across all the views. Notice that, depending on the view, the number of pixels mapped into the template may be smaller than the overall size of the template triangle. This might introduce a consistent number of missing pixels to the registered triangles. Fig. 2 shows the missing pixel values for different registered image patches as green points. After registering all the triangles \tilde{t}_μ from image frames $1 \dots f$ to the corresponding template, we can form the multi-view intensity matrix such as:




Figure 2: Multi-view patches in different frames are projected to the corresponding template creating a stack of images where pixel correspondences are given. Green pixels show the missing values after registering multi-view image patches to the template triangle.

$$\mathbf{J}_\mu = \begin{bmatrix} \text{vec}(\bar{\mathbf{J}}_{1\mu})^\top \\ \vdots \\ \text{vec}(\bar{\mathbf{J}}_{f\mu})^\top \end{bmatrix}_{f \times b}, \quad (9)$$

where vector $\text{vec}(\bar{\mathbf{J}}_{g\mu})$ ($g = 1 \dots f$) is the column-wise vectorization of registered image patch $\bar{\mathbf{J}}_{g\mu}$. Note that, registered image patches, i.e., $\bar{\mathbf{J}}_{g\mu}$, are $h \times w$ matrices with $|\mathcal{K}_\mu|$ non-zero values, where \mathcal{K}_μ is the set of pixel indices in $\bar{\mathbf{J}}_{g\mu}$ that have non-zero values and represent the registered triangle \hat{t}_μ (similar to Eq. (3)).

2.4 Surface Patch Reconstruction

After registering all the corresponding multi-view image patches, it is possible to compute the photometric parameters (albedo, surface normals, and lighting directions) from the matrix \mathbf{J}_μ for each multi-view image patch. A grayscale pixel $x_{gi}^{(\mu)} \in \mathbf{J}_\mu$ (pixel i from frame g belonging to triangular patch μ) can be defined in terms of a normal to the surface $\mathbf{n}_{\mu i}$, the albedo $\rho_{\mu i}$ at that pixel position and the lighting direction $\mathbf{l}_{\mu g}$ such that:

$$x_{gi}^{(\mu)} = \mathbf{l}_{\mu g}^\top \rho_{\mu i} \begin{bmatrix} 1 \\ \mathbf{n}_{\mu i} \end{bmatrix} \quad (10)$$

where $\mathbf{l}_{\mu g} \in \mathbb{R}^4$, $\rho_{\mu i} \in \mathbb{R}$, $\mathbf{n}_{\mu i} \in \mathbb{R}^3$ with the non-linear constraint $\mathbf{n}_{\mu i}^\top \mathbf{n}_{\mu i} = 1$. Rewriting Eq. (10) for all the pixels in patch μ and all the frames, according to [12] yields:

$$\mathbf{J}_\mu = \begin{bmatrix} \mathbf{l}_{\mu 1}^\top \\ \vdots \\ \mathbf{l}_{\mu f}^\top \end{bmatrix} \cdot \begin{bmatrix} \begin{bmatrix} \rho_{\mu 1} \\ \mathbf{n}_{\mu 1} \end{bmatrix} & \dots & \begin{bmatrix} \rho_{\mu b} \\ \mathbf{n}_{\mu b} \end{bmatrix} \end{bmatrix} = \mathbf{L}_\mu \cdot \mathbf{N}_\mu, \quad (11)$$

where the $f \times 4$ matrix \mathbf{L}_μ contains the collection of lighting directions for f frames, while the $4 \times b$ matrix \mathbf{N}_μ contains the surface normal vectors and albedo values.

However, the previous multi-view registration step has produced a matrix \mathbf{J}_μ with missing values and thus any closed-form solutions such as [12] cannot be applied to reconstructing the surface patch. To deal with this kind of missing data, we first define a binary mask for each patch μ as a matrix \mathbf{D}_μ , containing zeros for the elements corresponding to missing values. As an additional source of missing data, we also consider extremely dark and saturated pixels, since they do not provide relevant information about the photometric properties of the shape. Thus, the estimation problem can be formalized as an optimization that solves for photometric properties of a single patch μ such that:

$$\begin{aligned} \min_{\mathbf{L}_\mu, \mathbf{N}_\mu} \quad & \|\mathbf{D}_\mu \odot (\mathbf{J}_\mu - \mathbf{L}_\mu \mathbf{N}_\mu)\|^2, \\ \text{subject to} \quad & \mathbf{N}_\mu \in \mathcal{N} \end{aligned} \quad (12)$$

where

$$\mathcal{N} = \{\rho [1 \ \mathbf{n}^\top]^\top : \rho \in \mathbb{R}, \ \mathbf{n} \in \mathbb{R}^3, \ \mathbf{n}^\top \mathbf{n} = 1\}$$

is the *photometric manifold* representing the non-linear constraint given by the 1st order spherical harmonic approximation, as in Eq. (10). This bilinear optimization problem with manifold constraints can be solved using a general purpose solver, such as BALM [26]. Notice that we solve it for the surface normals associated to each pixel but not for the overall 3D surface. Therefore, a further integration step is required to recover the final 3D surface from the surface normals [27]. This step provides the 3-vector $\tilde{\mathbf{s}}_{\mu i}$ related to pixel $x_i^{(\mu)}$ that represents the 3D surface point i belonging to triangle μ .

Considering all triangular patches, m small PS problems have to be solved. This piecewise formulation for photometric stereo entails several advantages with respect to previous global approaches. Each patch has been associated with an individual photometric model rather than a single global model, as used in classical approaches. In this way, the piecewise formulation may grasp more complex lighting effects. Moreover, it is intrinsically more efficient and highly parallelizable, since reconstructing the surface in patches is computationally faster than reconstructing the global surface.

The next Section presents the final step of the pipeline and shows how to re-assemble the local patches to a global reference system given the coarse 3D mesh.

2.5 Surface Patch Alignment

A global and dense surface can be obtained only by registering each reconstructed surface patch to the corresponding position on the coarse 3D mesh \mathcal{S} . Such mapping is possible through the inverse of transformations \mathbf{R}_μ in Eq. (6) to align all surface patches back to the 3D mesh. In this way, we can map image pixels belonging to the registered triangle patch $\tilde{\mathbf{J}}_\mu$ to the corresponding facet of the 3D mesh \mathcal{S} as below:

$$\tilde{\mathbf{P}}_\mu = \frac{1}{\zeta} \mathbf{R}_\mu^{-1} \begin{bmatrix} \mathbf{x}_\mu^\top \\ \mathbf{y}_\mu^\top \\ 1 \end{bmatrix} \quad (13)$$

where vectors \mathbf{x}_μ and \mathbf{y}_μ include x and y components of the pixels belonging to registered image patch $\tilde{\mathbf{J}}_\mu$, and $\tilde{\mathbf{P}}_\mu$ is a matrix containing points of the corresponding facet of the 3D mesh.

Note that the transformed triangular plane is larger than the corresponding facet since we map the projection of enlarged triangles to the 3D mesh. These additional pixels help us to find the photo-geometric transformation since they are overlapping points along the borders of proximate triangular patches. These overlapping points should be aligned with the corresponding ones on the neighbouring surface patches. This constraint is used to obtain the final surface. We can represent the surface patch $\tilde{\mathbf{S}}_\mu$ on the 3D mesh as transformed image points from triangular image patch $\tilde{\mathbf{J}}_\mu$ that are elevated along the facet normal \mathbf{N}_f^μ with respect to the recovered surface patch $\tilde{\mathbf{S}}_\mu$, where

$$\tilde{\mathbf{S}}_\mu = [\tilde{\mathbf{s}}_{\mu 1} \dots \tilde{\mathbf{s}}_{\mu |\mathcal{K}_\mu|}]_{3 \times |\mathcal{K}_\mu|}, \quad (14)$$

accumulates $|\mathcal{K}_\mu|$ 3D vectors representing points on the surface of triangular patch \tilde{t}_μ .

Therefore, the relation between the overlapping points for surface patches belonging to neighboring triangles \tilde{t}_μ and $\tilde{t}_{\mu'}$ can be described considering that the corresponding points from two surface patches should be identical on the global surface \mathbf{S} . This relation is represented as:

$$\begin{aligned} \tilde{\mathbf{P}}_\mu^{(c)} + \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^\mu \end{bmatrix}_{3 \times 3} \cdot \mathbf{H}_\mu \cdot \tilde{\mathbf{S}}_\mu^{(c)} &= \\ \tilde{\mathbf{P}}_{\mu'}^{(c)} + \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^{\mu'} \end{bmatrix}_{3 \times 3} \cdot \mathbf{H}_{\mu'} \cdot \tilde{\mathbf{S}}_{\mu'}^{(c)}, & \end{aligned} \quad (15)$$

where $\tilde{\mathbf{P}}_\mu^{(c)}$ represents the common points between $\tilde{\mathbf{P}}_\mu$ and $\tilde{\mathbf{P}}_{\mu'}$, $\tilde{\mathbf{S}}_\mu^{(c)}$ give the corresponding points between $\tilde{\mathbf{S}}_\mu$ and $\tilde{\mathbf{S}}_{\mu'}$ and $\mathbf{0}_3$ is a vector of 3 zeros. We can rewrite Eq. (15) as:

$$\begin{aligned} \text{vec}(\tilde{\mathbf{P}}_\mu^{(c)}) + \left(\left(\tilde{\mathbf{S}}_\mu^{(c)} \right)^\top \otimes \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^\mu \end{bmatrix} \right) \cdot \text{vec}(\mathbf{H}_\mu) = \\ \text{vec}(\tilde{\mathbf{P}}_{\mu'}^{(c)}) + \left(\left(\tilde{\mathbf{S}}_{\mu'}^{(c)} \right)^\top \otimes \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^{\mu'} \end{bmatrix} \right) \cdot \text{vec}(\mathbf{H}_{\mu'}), \end{aligned} \quad (16)$$

where operator \otimes is the *Kronecker product*.

Writing Eq. (16) for all the points overlapping between every pair of triangles, the problem of estimating matrices \mathbf{H}_μ and $\mathbf{H}_{\mu'}$ can be formalized as an

optimization problem:

$$\min_{\mathbf{H}_\mu, \mathbf{H}_{\mu'}} \left\| \begin{pmatrix} \left(\left(\tilde{\mathbf{S}}_\mu^{(c)} \right)^\top \otimes \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^\mu \end{bmatrix} \right) \cdot \text{vec}(\mathbf{H}_\mu) - \\ \left(\left(\tilde{\mathbf{S}}_{\mu'}^{(c)} \right)^\top \otimes \begin{bmatrix} \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{N}_f^{\mu'} \end{bmatrix} \right) \cdot \text{vec}(\mathbf{H}_{\mu'}) + \\ \text{vec}(\tilde{\mathbf{P}}_\mu^{(c)}) - \text{vec}(\tilde{\mathbf{P}}_{\mu'}^{(c)}) + k \end{pmatrix} \right\|^2, \quad (17)$$

$$k = \frac{1}{\sum(C_\mu + C_{\mu'})}$$

where C_μ and $C_{\mu'}$ denote the surface curvatures for triangles \tilde{t}_μ and $\tilde{t}_{\mu'}$. The curvature term k prevents the solutions for \mathbf{H}_μ and $\mathbf{H}_{\mu'}$ from being chosen in a way that flattens the surface patches $\tilde{\mathbf{S}}_\mu$. Solving Eq. (17) for the transformations \mathbf{H}_μ and $\mathbf{H}_{\mu'}$ with Least Squares aligns all the surface patches by solving their local bas-relief ambiguity up to an arbitrary reduced local ambiguity.

The global dense surface $\hat{\mathcal{S}}$ is obtained after variational refinement of the surface through minimizing the energy function:

$$E(\hat{\mathcal{S}}) = \int \lambda(1 - G(\mathbf{s})) \left\| \bigcup_\mu \mathbf{S}_\mu - \hat{\mathcal{S}} \right\|_2^2 + G(\mathbf{s}) \left\| \nabla \hat{\mathcal{S}} \right\|_2^2 \text{d}\mathbf{s}, \quad (18)$$

where \mathbf{s} is a surface point, $\bigcup \mathbf{S}_\mu$ is the geometric superposition of surface patches, and $\lambda > 0$ is a parameter to control the smoothness. $G(\cdot)$ is the weighting function

$$G(\mathbf{s}) = \frac{1}{2} \sqrt{\|\alpha^{\mathbf{s}} - \beta^{\mathbf{s}}\|_2 + \|\alpha^{\mathbf{s}} - \gamma^{\mathbf{s}}\|_2 + \|\gamma^{\mathbf{s}} - \beta^{\mathbf{s}}\|_2}, \quad (19)$$

where $\alpha^{\mathbf{s}}$, $\beta^{\mathbf{s}}$ and $\gamma^{\mathbf{s}}$ are the barycentric coordinates of 3D point \mathbf{s} with respect to the corresponding triangle.

3 Experiments

The proposed 3D reconstruction pipeline was tested on several sequences from *RobotDataset* [28]¹. The main feature of this dataset is the possibility to have a dense 3D ground truth obtained by structured light. This represents a step forward with respect to the previous evaluation protocols where experimental results were mostly assessed on a qualitative basis. In terms of PS reconstruction, the sequences offer several challenges: i) the setup uses LED lights for illumination and it is completely uncalibrated; ii) strong shadows affect the images; iii) the general object's surface properties depart from a Lambertian model, iv) the corresponding PS problem is also large scale in its nature. Each single image has a frame size of 1600×1200 pixels, therefore, each image sequence has a number of elements (pixels) in the order of 10^9 .

¹<http://roboimagedata.imm.dtu.dk>

Comparison with previous approaches Given a test sequence in the dataset, we select 30 frames with both moving camera and varying light sources. The camera motion on the arc gives a variation of 50 degrees overall. In addition to the wide base-line, the images are high resolution and present strong shadows and reflection effects. Figure 3 shows the results for three sequences in the dataset (*Mug*, *Jar* and *Ball*) compared with a classical MVS (i.e., Furukawa and Ponce [29]), and the MVPS approach of Lim et al. [7].

The *Mug* sequence has a glossy surface with a high reflectance but it is highly textured. Such properties make it challenging for standard PS but the texture is an advantage for MVS methods. In particular, this sequence shows that our approach can provide results comparable with MVS methods. The *Ball* sequence has strong cast shadows and complex bending, which introduce interesting challenges for MVS methods. Although the surface is not reflective and can be considered as being almost Lambertian, the extreme lighting conditions give also a challenge to standard PS approaches. The *Jar* sequence also has a glossy surface with many saturated pixels which make the photometric reconstruction very challenging. Moreover, the smooth convexity and lack of sufficient texture introduce difficulties for MVS methods.

Multi-view photometric stereo approaches have the disadvantage of requiring short baseline images. As a practical example, we have tested the method by Lim et al. [7], which has similarities with our approach. We have noticed that the main problem affecting this approach is the difficulty in finding correct pixel-to-pixel correspondences when the baseline between views increases. In most of the tested cases, the recovered shape by Lim et al [7] is almost flat with some bumps and few spikes (Figure 3 - the 3rd column).

In contrast, multi-view stereo approaches can deal with wide baseline images but the reconstruction is not dense enough to include details of the surface. The right column of Figure 3 (the 4th column) shows the reconstruction from Furukawa and Ponce [29]. Although the overall geometry of the recovered shape is acceptable, the surface can only be recovered for highly-textured parts of the shape, like in *Mug* sequence. The strong shadows in the *Ball* sequence prevented this method from matching pixels on some parts of the surface, which is also the case in *Jar* sequence because of poor texture.

Computational gain Our piecewise formulation provides an efficient solution in terms of computational requests. This comes directly from considering several smaller PS problems instead of a single large-scale problem. In general, a piecewise formulation deals with multi-view patches of about 10^4 pixels while a single PS example might amount to 10^8 pixels approximately. To show such an effect, we compared the processing time and the performance for the *potato* sequence example in [26] using static view PS and an analogous sequence with our MVPS. The reported running time [26] is almost 22 hours and resulted in 17.34% for mean 3D errors, instead our method achieved 11.15% 3D error in 50.8 minutes.

Table 1 provides some quantitative data regarding our solution. The rate of




Figure 3: Images from *Mug*, *Jar* and *Ball* sequences together with our results (error heatmaps and 3D surface in the 1st and 2nd columns) and comparisons with a MVPS method [7] (3rd column) and a MVS method [29] (4th column).

missing pixel values is noticeable and the obtained mean 3D error shows that the PPS could handle remarkable amounts of missing data. The timing presented in this table includes reconstruction of all the patches taking into account for the photo-geometric alignment.

Table 1: Quantitative data regarding our solution for different sequences.

Sequence	# Mesh Points	# Triangles	# Pixels	% Missing Pixels	Time (Min.)	3D Error
Mug	115	198	6,689,940	24.7%	24.38	7.1%
Jar	130	206	22,618,080	57.79%	33.14	9.34%
Balls	250	362	42,374,490	69.48%	62.63	9.69%
Potatoes	150	276	36,672,990	62.72%	50.8	11.15%

4 Conclusions

We have presented a novel photo-geometric method for dense reconstruction from multi-view with arbitrary lighting condition. The approach is able to cope with wide-baselines images from uncalibrated cameras and explicitly utilizes varying lighting directions in the image sequence. This means that, only a few images taken by the end user are enough to be able to recover the dense 3D surfaces. The piecewise approach is highly scalable since solving for image patches is computationally easier than considering whole images. This enables the pipeline to run on commodity PCs. Future work will be dedicated to studying approaches that can partition the image into a mesh while taking into account the photometric and geometric properties of the shape. For instance, the method in [30] could be used to partition the image in more consistent patches allocated to different subspaces. In addition, more complex photometric models could be used to extract more realistic photometric attributes for the surface.

Acknowledgement

This work is supported by the Doctoral fellowship awarded by the Italian Government and the Italian Institute of Technology. The authors would like to thank Prof. Davide Scaramuzza for inspiring discussions and his valuable comments.

References

- [1] S. Agarwal, N. Snavely, S. M. Seitz, and R. Szeliski, “Bundle adjustment in the large,” in *European Conference on Computer Vision (ECCV’10)*, pp. 29–42, Springer, 2010.
- [2] R. A. Newcombe and A. J. Davison, “Live dense reconstruction with a single moving camera,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR’10)*, pp. 1498–1505, IEEE, 2010.

- [3] K. Matzen and N. Snavely, “Scene chronology,” in *European Conference on Computer Vision (ECCV’14)*, pp. 615–630, Springer, 2014.
- [4] M. Pizzoli, C. Forster, and D. Scaramuzza, “Remode: Probabilistic, monocular dense reconstruction in real time,” in *IEEE International Conference on Robotics and Automation (ICRA’14)*, pp. 2609–2616, IEEE, 2014.
- [5] J. Engel, T. Schöps, and D. Cremers, “Lsd-slam: Large-scale direct monocular slam,” in *European Conference on Computer Vision (ECCV’14)*, pp. 834–849, Springer, 2014.
- [6] L. Zhang, B. Curless, A. Hertzmann, and S. M. Seitz, “Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multiview stereo,” in *IEEE International Conference on Computer Vision (ICCV’03)*, pp. 618–625, IEEE, 2003.
- [7] J. Lim, J. Ho, M.-H. Yang, and D. Kriegman, “Passive photometric stereo from motion,” in *IEEE International Conference on Computer Vision (ICCV’05)*, vol. 2, pp. 1635–1642, IEEE, 2005.
- [8] N. Joshi and D. Kriegman, “Shape from varying illumination and viewpoint,” in *IEEE International Conference on Computer Vision (ICCV’07)*, pp. 1–7, IEEE, 2007.
- [9] C. Hernández, G. Vogiatzis, and R. Cipolla, “Multiview photometric stereo,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI’08)*, vol. 30, no. 3, pp. 548–554, 2008.
- [10] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, “Building rome in a day,” *Communications of the ACM*, vol. 54, no. 10, pp. 105–112, 2011.
- [11] H. Hayakawa, “Photometric stereo under a light source with arbitrary motion,” *Journal of Optical Society of America*, pp. 3079–3089, 1994.
- [12] R. Basri, D. Jacobs, and I. Kemelmacher, “Photometric stereo with general, unknown lighting,” *International Journal of Computer Vision (IJCV’07)*, vol. 72, no. 3, pp. 239–257, 2007.
- [13] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille, “The bas-relief ambiguity,” *International Journal of Computer Vision (IJCV’99)*, vol. 35, no. 1, pp. 33–44, 1999.
- [14] B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan, “Self-calibrating photometric stereo,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR’10)*, pp. 1118–1125, IEEE, 2010.
- [15] T. Papadhimetri and P. Favaro, “A closed-form, consistent and robust solution to uncalibrated photometric stereo via local diffuse reflectance maxima,” *International Journal on Computer Vision (IJCV’13)*, pp. 1–16, 2013.

- [16] D. Scharstein and R. Szeliski, “High-accuracy stereo depth maps using structured light,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’03)*, vol. 1, pp. I–195, IEEE, 2003.
- [17] J. Salvi, J. Pages, and J. Batlle, “Pattern codification strategies in structured light systems,” *Pattern Recognition*, pp. 827–849, 2004.
- [18] C. Hernández, G. Vogiatzis, G. J. Brostow, B. Stenger, and R. Cipolla, “Non-rigid photometric stereo with colored lights,” in *IEEE International Conference on Computer Vision (ICCV’07)*, pp. 1–8, IEEE, 2007.
- [19] R. Anderson, B. Stenger, and R. Cipolla, “Color photometric stereo for multicolored surfaces,” in *IEEE International Conference on Computer Vision (ICCV’11)*, pp. 2182–2189, IEEE, 2011.
- [20] J. Park, S. N. Sinha, Y. Matsushita, Y.-W. Tai, and I. S. Kweon, “Multiview photometric stereo using planar mesh parameterization,” in *IEEE International Conference on Computer Vision (ICCV’13)*, pp. 1161–1168, IEEE, 2013.
- [21] A. Varol, M. Salzmann, E. Tola, and P. Fua, “Template-free monocular reconstruction of deformable surfaces,” in *IEEE International Conference on Computer Vision (ICCV’09)*, pp. 1811–1818, IEEE, 2009.
- [22] J. Taylor, A. D. Jepson, and K. N. Kutulakos, “Non-rigid structure from locally-rigid motion,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR’10)*, pp. 2761–2768, IEEE, 2010.
- [23] J. Fayad, L. Agapito, and A. Del Bue, “Piecewise quadratic reconstruction of non-rigid surfaces from monocular sequences,” in *European Conference on Computer Vision (ECCV’10)*, pp. 297–310, Springer, 2010.
- [24] R. Sabzevari, A. Del Bue, and V. Murino, “Piecewise single view photometric stereo with multi-view constraints,” in *International Conference on Image Processing (ICIP’12)*, pp. 21–24, 2012.
- [25] B. Horn, “Closed-form solution of absolute orientation using unit quaternions,” *Journal of Optical Society of America*, pp. 629–642, 1987.
- [26] A. Del Bue, J. Xavier, L. Agapito, and M. Paladini, “Bilinear modeling via augmented lagrange multipliers (balm),” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI’12)*, vol. 34, no. 8, pp. 1496–1508, 2012.
- [27] R. T. Frankot and R. Chellappa, “A method for enforcing integrability in shape from shading algorithms,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI’88)*, vol. 10, no. 4, pp. 439–451, 1988.
- [28] H. Aanæs, A. L. Dahl, and K. S. Pedersen, “Interesting interest points,” *International Journal of Computer Vision (IJCV’11)*, vol. 97, no. 1, pp. 18–35, 2012.

- [29] Y. Furukawa and J. Ponce, “Accurate, dense, and robust multiview stereopsis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI’10)*, pp. 1362–1376, 2010.
- [30] K. Sunkavalli, T. Zickler, and H. Pfister, “Visibility subspaces: Uncalibrated photometric stereo with shadows,” in *European Conference on Computer Vision (ECCV’10)*, pp. 251–264, Springer, 2010.